

This document contains a list of comments Will MacAskill has made about Sam Bankman-Fried before November 11th, 2022; it's been put together by a [Forethought](#) employee. The scope is anything that was publicly said or has since been made public. Inevitably, this search won't be wholly exhaustive and some quotes will be missing, but significant effort went into being as thorough as possible. If you find a quote within the scope of this document not listed below, please email contact@forethought.org with a link to the quote. Quotes were sourced through:

- I. Google searches for various combinations and alterations of “*macaskill*” AND “*fried*” before: 2023-11-12” and clicking through dozens of articles searching for quotes.
- II. Will MacAskill's twitter feed (including replies)
- III. Will MacAskill's EA forum account (including comments)
- IV. Podcasts Will MacAskill has gone on
- V. Interviews Will MacAskill has given
- VI. Talks Will MacAskill has given

[1. Atlantic: Effective Altruism's Philosopher King Just Wants to Be Practical](#)

[2. 80,000 Hours Podcast](#)

[3. New Yorker](#)

[4. EA and the current funding situation](#)

[5. Ali Abdaal's podcast](#)

[6. Elon Musk has a go-to philosopher. Meet William MacAskill](#)

[7. Retweeing Sam's tweet about being proud of the Future Fund](#)

[8. Retweeting Sam's thread on the announcement of Future Fund](#)

[9. William MacAskill - Caring About Future People Is Common Sense | The Daily Show](#)

[10. People I \(Mostly\) Admire:](#)

1. [Atlantic: Effective Altruism's Philosopher King Just Wants to Be Practical](#)

- a. Stern: Shifting out of pandemic gear, I was wondering whether there are major lobbying efforts under way to persuade billionaires to convert to EA, given that the potential payoff of persuading someone like Jeff Bezos to donate some significant part of his fortune is just massive.
- b. MacAskill: I do a bunch of this. I've spoken at the Giving Pledge annual retreat, and I do a bunch of other speaking. It's been pretty successful overall, insofar as there are

other people kind of coming in—not on the size of Sam Bankman-Fried or Dustin Moskovitz and Cari Tuna, but there’s definitely further interest, and it is something I’ll kind of keep trying to do. Another organization is Longview Philanthropy, which has done a lot of advising for new philanthropists to get them more involved and interested in EA ideas.

2. [80,000 Hours Podcast](#)

- a. Will MacAskill: For sure. It’s definitely true that not everyone can do that, so you’ve got to have some argument about why you’re in an unusual position. But I at least think for me that was just absolutely correct. **So I gave a talk at MIT about earning to give and Sam Bankman-Fried was there and he went and earned to give, and now he’s doing rather well. That was really the important impact.**
- b. Will MacAskill: So yeah, this theme is “A Culture of Ambition.” Why, if you’re altruistic, should you try and be very ambitious, or have ambitious plans? I think there’s a few reasons for that. One is more theoretical: if you think the distribution of impact is a fat-tailed distribution, that means that the mean impact is bigger than the median. So if you just aim to do kind of a typical amount of good, then you’re going to do much less good in expectation — and I think plausibly it’s the expected value of the good that you do that matters — than if you’re really aiming for the best. And that’s because, in the nature of the distribution, the best outcomes are just way, way better than the typical outcomes.

Will MacAskill: So there’s this theoretical case, and we’ve seen that in practice as well. I think it’s clearest if you look at people earning to give, because there, we can just assign a number on the impact they’re going to have — at least the impact in terms of donations — where the distribution of how much money you can make from careers that you might go into, earning to give, is just clearly this fat-tailed distribution.

Will MacAskill: When we look at the people who’ve earned the most in order to donate via earning to give, well, it’s enormously driven by those at FTX now. Dustin Moskovitz is obviously another major donor — my understanding is that he got into giving after cofounding Facebook; I don’t know about his motivations before that, whether he was always planning to give. **But certainly for Sam and the other early employees at FTX, their aim was to make money in order to donate it. And the amount they’ve raised now is just very large compared to how much they could have made, if Sam had stayed working at Jane Street, for example, even though Jane Street’s a very well-paying quantitative trading firm.**

Rob Wiblin: Yeah. Many listeners will have listened to the interview with Sam Bankman-Fried, but a couple won't have, so it might be worth saying something quickly about the kind of sums that we're talking about.

Will MacAskill: Yeah, that's right. **So Sam intends to give away essentially all of his wealth, like 99%, as do the other early employees, is my understanding. I don't know the details for each person. Sam's net worth, at the moment, was estimated by Forbes to be \$24 billion. Gary Wang, who has helped create FTX, his net worth is now public at about \$6 billion. And there are a few other early employees of FTX who are also doing it for earning-to-give reasons as well. So we are talking about very large sums of money now.**

Rob Wiblin: Yeah. Sorry, go on. What implications does that have?

Will MacAskill: So that's just an illustration in practice of this fact that impact can be fat tailed. We focused on earning to give — that's nice and quantifiable. But I'm sure this is true as well in other areas. Within politics, if you manage to become a head of state, then you're having very unusually large impact. If you're doing research, like being one of the seminal researchers, you're going to have enormous impact compared to if you are a typical researcher.

Will MacAskill: And that really means that if you want to do as much good in expectation as possible, ideally you should be aiming for those best-case outcomes. And probably, you'll fail — more likely than not. **When Sam set up FTX, he thought it was a 20% chance of succeeding — and he was being relatively optimistic, he thought at the time. But obviously, he did end up being successful. But even if you divide that by five, expected value —**

Rob Wiblin: It's still a big number. Yeah.

Will MacAskill: It's still bigger, and it's bigger than what he would've made if he'd stayed at Jane Street. Similarly, when people have looked into this — like, "Should you go into politics?" — obviously it's extremely hard to become an MP in the UK. And even harder again to become, say, the prime minister. But the chance of achieving one of these best-case outcomes is plausibly where much of the expected value comes from. And I think the same could be true for research and other career paths.

Rob Wiblin: So it seems like, over the last 10 years, we've had a decent number of people who've had this kind of "swing for the fences" approach to their careers or to doing good. Do you think we're still short of what would be optimal, in terms of how risk-taking people typically are?

Will MacAskill: I think we should expect that to be the case, where this is a way in which altruistic reasoning might be quite different from, say, self-interest. Again, it's

easiest thinking about earning to give: if you're just earning money for yourself, then the difference between making 100,000 and 200,000 is very great. The difference between making 1,000,000 a year and 1,100,000 a year is not nearly as great. Whereas altruistically, plausibly, those differences matter basically about as much. And that's quite an unusual thing to kind of think about.

- c. Will MacAskill: Exactly. You've got much larger datasets on gang violence. Whereas if you're concerned about a war between the US and China, probably just details of the particular situation are going to matter a lot more than a dataset that includes several conflicts from the 1900s or something.

Rob Wiblin: Yeah, when I spoke with Sam, he was really advocating for this approach to evaluating the grants where you spend one minute reading the grant, and then you say, "If I spent another week evaluating this topic, how likely would it be to change my decision?" And if the answer is it probably wouldn't change the decision, then you just go, "Yes" or "No" and then move on. I suppose with 1,800 applications, you kind of don't have much choice but to adopt that.

Will MacAskill: **Well, bearing in mind that 90% of them were easy noes. But there was one part of that conversation with Sam that really resonated with me, where you've got a group of people discussing something, and everyone thinks it should be funded, but some people "have concerns" and they just want to talk about their concerns. But even after the concerns, they still want to fund it, just like, "Yeah, but I think we should just feel a bit bad about that," or something. That's a mode I think is easy to get into and we should try and avoid.**

- d. Rob Wiblin: A final one on the Future Fund. Do you have any interesting, exciting disagreements with Sam Bankman-Fried or Nick Beckstead?

Will MacAskill: It's a great question. **I actually find myself remarkably aligned with Sam, in interesting ways, on a couple of issues. One is this sympathy to broad longtermism. I think the two issues relate to each other quite closely. One worldview you might have is that there are things that are narrowly focused on extinction risk, and those things have longer-term impact, and basically nothing else does. And then you've got this different worldview, which Sam was espousing on this podcast, which is like, the world's really connected — loads of things you can do have indirect impact on things that then do impact the very long term.**

Will MacAskill: I'm much more in the latter camp. And hopefully the book gives some kind of background explanation for that. But part of that is this kind of contingency

versus convergence-of-values idea as well. Where if you think it really matters who's in charge, who has the power at these critical junctures — like development of AGI or first space exploration or formation of a world government — then loads of things change that. And that means you can affect the long-term future in just an enormous number of ways.

Will MacAskill: **Probably the biggest disagreement I'd have with Sam is like, it's kind of a spectrum of how cautious versus gung-ho you are. Sam is very keen on building things, just making things happen. And I'm really pretty sympathetic to that compared to many people, but he's definitely further on the tail of that than me.**

3. [New Yorker](#)

- a. **Last year, MacAskill contacted Bankman-Fried to check in about his promise: “Someone gets very rich and, it's, like, O.K., remember the altruism side? I called him and said, ‘So, still planning to donate?’”** Bankman-Fried pledged to give nearly all his money away; if suitable opportunities are found, he's willing to contribute more than a billion dollars a year. **Bankman-Fried had longtermist views before they held sway over MacAskill, and has always been, MacAskill remembers, “particularly excited by pandemics”**—a normal thing to hear among E.A.s. Bankman-Fried set up a foundation, the FTX Future Fund, and hired the longtermist philosopher Nick Beckstead as C.E.O. This past December, MacAskill finished the manuscript of his new book, and hoped to spend more time with his partner, Holly Morgan, an early E.A. and the biggest single input to his stability. Instead, Bankman-Fried enlisted him as a Future Fund adviser. (He offered MacAskill a “generous” six-figure salary, but MacAskill replied that he was just going to redistribute the money anyway.)

4. [EA and the current funding situation](#)

- a. The primary update comes from the success of FTX: Sam Bankman-Fried has an estimated net worth of \$24 billion (though bear in mind the difficulty of valuing crypto assets, and their volatility), and intends to give essentially all of it away. The other EA-aligned FTX early employees add considerably to that total.
- b. **Heavily considering what you show as well as what you do, especially if you're in a position of high visibility. “Signalling” is often very important! For example, the funding situation means I now take my personal giving more seriously, not less. I think the fact that Sam Bankman-Fried is a vegan and**

drives a Corolla is awesome, and totally the right call. And, even though it won't be the right choice for most of us, we can still celebrate those people who do make very intense moral commitments, like the many kidney donors in EA, or a Christian EA I heard about recently who lives in a van on the campus of the tech company he works for, giving away everything above \$3000 per year.

- c. I worry we've made some errors in the past by not taking the chance of best-case scenarios seriously, out of a desire to be rigorous and skeptical. For example, I mentioned that Toby initially estimated the value of a Giving What We Can Pledge at \$70,000 (as one example of quantifying the benefits of outreach and community-building more generally). I remember having arguments with people who claimed that estimate was too optimistic. But take the 7000 Giving What We Can members, and assume that none of them give anything apart from Sam Bankman-Fried, who gives his net worth. Then a Pledge was actually worth \$2 million — 30 times higher than Toby's "optimistic" estimate at the time.[16] In general, if our successes are sampling from a heavy-tailed distribution, the historical average value of our impact will very likely be lower than the true mean.
5. [Ali Abdaal's podcast](#)
 - a. WM, when asked why 80,000 hours do advertisements on podcasts: Let's just take one example of impact which was in 2012, I was giving this talk about earning to give and one of the talks was at MIT and there was an audience member there who was called sam who was like "oh these ideas make sense" and he went to earn to give and did quantitative trading at Jane Street and then moved on to set up his own company. I believe he is now the richest person in the world under the age of 35, has already given hundreds of millions of dollars through his foundation this year, and has publicly said he will give away ya know essentially all of his wealth, 99% of his wealth.
 6. [Elon Musk has a go-to philosopher. Meet William MacAskill](#)
 - a. Bankman-Fried and MacAskill go back a long way together. He was still an undergraduate at MIT when he met MacAskill. **The two men kept in touch and as Bankman-Fried's wealth grew, he became a major figure in the effective altruism movement, serving on the board of the Centre for Effective Altruism and, as MacAskill proudly remarked, was on course to "donate more than anyone in history before the age of 35". This extraordinary generosity "was in significant part on the basis of meeting me".** When Elon Musk's phone was hacked, messages were discovered on it from MacAskill suggesting he put Musk in

touch with Bankman-Fried with a view to a “joint effort” between the two billionaires to buy Twitter. Nothing came of it.

7. [Retweeing Sam’s tweet about being proud of the Future Fund](#)
8. [Retweeting Sam’s thread on the announcement of Future Fund](#)
9. [William MacAskill - Caring About Future People Is Common Sense | The Daily Show](#)
 - a. WM: So, I think it’s absolutely an obligation on the most privileged. So, yeah, if you’re a billionaire. So, Bill Gates, you know, ushered in a new era of billionaire philanthropy, saying 50%.
 - b. TN: Right
 - c. WM: I think the standard should be higher: 99% or more. And ...
 - d. TN: *laughing* *audience applause and laughing*
 - e. WM: And ... uh, so, yeah, um, people are laughing, but a decent chunk of my life is talking to billionaires and trying to get them - convince them to give almost all their money away.
 - f. TN: What do you say? I want to know what you say, William. Help me understand this. Help me understand.
 - g. WM: So, I mean, uh, so, yeah.
 - h. TN: Cause billionaires are notorious for not giving their money away.
 - i. WM: I think that’s gonna change. Um, at least -
 - j. TN: So, what do you say? What do you say to them?
 - k. WM: So, at least for some. So, say two things. One, like you say, just, money just, above a certain threshold, money just doesn’t affect your well-being, come on.
 - l. TN: Right
 - m. WM: And then, secondly money can be used to make a transformative difference. It costs \$3,000 to save a life by protecting children in sub-Saharan Africa, uh, from malaria by distributing insecticide treated bed nets. This is a very reliable, well-evidenced fact.
 - n. ...
 - o. WM: The case is just like, ‘look, it’s gonna make almost no difference in your life. It’s gonna be enormously impactful for the world.
 - p. TN: Right
 - q. WM: ‘And actually, we have seen success. So, there are a number of people, um, in the kind of broader effective altruism movement Dustin Moskowitz, now Sam

Bankman-Fried who are planning to give 99% or more and I think that can be the new standard.

10. [People I \(Mostly\) Admire:](#)

- a. LEVITT: Now, my own experience is once you put yourself in an environment like investment banking, it's really hard to stay true to whatever fundamental beliefs you had when you started. And I guess that's an empirical question. It would be an interesting study to look into — to see how often people who say they're going to do good end up doing good and how many get sucked down the vortex.

MACASKILL: I think there's a small amount of data from effective altruists pursuing this path. And I think they might be a very different sample than the normal person who works in investment banking and claims they're going to give money away. I'm not sure how truly sincere those people were at the time. And this is certainly a worry I had. I called it the "corruption worry." But it seems to be not holding up, at least for the people who are really engaged in effective altruism. Those people I know, the highest earners are generally giving much more than 50 percent. **The kind of biggest success story for earning to give is Sam Bankman-Fried, who's now the richest person in the world under the age of 35, or at least that was true recently. And he's publicly stated he's giving away 99 percent of his wealth or more, and is already ramping up his giving.**

LEVITT: And he's a crypto guy?

MACASKILL: Yeah, he co-founded a crypto exchange called F.T.X. I think part of what helps here is that we've just built this community. It's much easier to live up to your ideals if you've got a bunch of people around you who will praise you for living up to them, such as giving. And maybe you'll feel shunned or less welcome if you're claiming that you want to do enormous amounts of good and give enormous amounts, but actually aren't. So, that's the kind of incentives, a count of perhaps how we've done better than the base rate of people working in investment banking.